

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re application of: **Bardsley et al.**

Serial No. **09/917,368**

Filed: **July 27, 2001**

For: **Correlating Network Information
and Intrusion Information to Find the
Entry Point of an Attack Upon a
Protected Computer**

§
§
§
§
§
§
§

Group Art Unit: **2137**

Examiner: **Jeffrey D. Popham**

**Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450**

37945
PATENT TRADEMARK OFFICE
CUSTOMER NUMBER

REPLACEMENT APPEAL BRIEF (37 C.F.R. 41.37)

This brief is in furtherance of the Notice of Appeal, filed in this case on September 27, 2006.

No fees are believed to be required. If, however, any fees are required, I authorize the Commissioner to charge these fees which may be required to IBM Corporation Deposit Account No. 09-0457. No extension of time is believed to be necessary. If, however, an extension of time is required, the extension is requested, and I authorize the Commissioner to charge any fees for this extension to IBM Corporation Deposit Account No. 09-0457.

REAL PARTY IN INTEREST

The real party in interest in this appeal is the following party: International Business Machines Corporation of Armonk, New York.

RELATED APPEALS AND INTERFERENCES

With respect to other appeals or interferences that will directly affect, or be directly affected by, or have a bearing on the Board's decision in the pending appeal, there are no such appeals or interferences.

STATUS OF CLAIMS

A. TOTAL NUMBER OF CLAIMS IN APPLICATION

Claims in the application are: 5-11 and 15-27.

B. STATUS OF ALL THE CLAIMS IN APPLICATION

1. Claims canceled: 1-4 and 12-14.
2. Claims withdrawn from consideration but not canceled: None.
3. Claims pending: 5-11 and 15-27.
4. Claims allowed: None.
5. Claims rejected: 5-11 and 15-27.
6. Claims objected to: None.

C. CLAIMS ON APPEAL

The claims on appeal are: 5-11 and 15-27.

STATUS OF AMENDMENTS

No amendments were submitted after the final office action of June 27, 2006.

SUMMARY OF CLAIMED SUBJECT MATTER

A. CLAIM 5 - INDEPENDENT

The subject matter of claim 5 deals with identifying the entry point of an attack upon a device protected by an intrusion detection system. See Specification, p. 1, ll. 4-6 and Figure 5.

The method of claim 5 includes the steps of obtaining intrusion information, from an intrusion detection system, regarding an attack upon a device protected by the intrusion detection system (Specification, p. 8, ll. 3-10; and Figure 2, reference numeral 200), obtaining network information, from network equipment connected to the device, regarding the attack (Specification, p. 8, l. 18 through p. 9., l. 10; and Figure 3, reference numeral 300), determining a logical entry point of the attack using a correlation engine to correlate the intrusion information and the network information (Specification, p. 6, l. 14 though p. 7, l. 2) and identifying a physical entry point associated with the logical entry point (Specification, p. 11, ll. 7-10; and Figure 5, reference numeral 530).

B. CLAIM 21 - INDEPENDENT

The subject matter of claim 21 deals with a computer-implemented method of identifying the entry point of an attack upon a device protected by an intrusion detection system, wherein the device is one of a plurality of devices connected by a network. See the specification, p. 1, ll. 4-6; specification p. 7, ll. 3-10; and Figure 5.

The method includes the computer-implemented steps of detecting an attack on the device (Specification, p. 7, ll. 3-10), notifying a correlation engine of the attack on the device (Specification, p. 10, ll. 6-12; and Figure 5, reference numeral 505), obtaining intrusion information regarding the attack (Specification, p. 8, ll. 3-10; and Figure 2, reference numeral 200), obtaining network information regarding the attack (Specification, p. 8, l. 18 through p. 9., l. 10; and Figure 3, reference numeral 300), using the correlation engine, correlating the intrusion information and the network information to produce correlation information (Specification, p. 6, l. 14 though p. 7, l. 2), using the correlation information, finding on the network a logical port of connection used by

the attack (Specification, p. 6, l. 14 though p. 7, l. 2), and mapping the logical port on the network to a physical port on the network using the correlation engine (Specification, p. 11, ll. 7-10; p. 9, ll. 15-20; and Figure 5, reference numeral 530).

C. CLAIM 25 - INDEPENDENT

The subject matter of claim 25 deals with an apparatus for detecting a point of an attack on a network. See the specification, p. 1, ll. 4-6; specification p. 7, ll. 3-10; and Figure 5.

The apparatus includes network equipment for connecting a protected device to a network (Specification, p. 7, l. 11 through p. 8, l. 2; and Figure 1, reference numeral 110), an intrusion detection system comprising intrusion detection equipment (Specification, p. 7, ll. 3-10), a correlation engine (Specification, p. 7, ll. 15-20; and Figure 1, reference numeral 140) adapted to receive a notification of an attack on the protected device (Specification, p. 10, ll. 6-12; and Figure 5, reference numeral 505), receive intrusion information regarding the attack (Specification, p. 8, ll. 3-10; and Figure 2, reference numeral 200), and receive network information regarding the attack, wherein the network information pertains to the network (Specification, p. 8, l. 18 through p. 9, l. 10; and Figure 3, reference numeral 300), correlate the intrusion information and the network information to produce correlation information (Specification, p. 6, l. 14 though p. 7, l. 2), use the correlation information to find on the network a logical port of connection used by the attack (Specification, p. 6, l. 14 though p. 7, l. 2), and map the logical port on the network to a physical port on the network using the correlation engine (Specification, p. 11, ll. 7-10; p. 9, ll. 15-20; and Figure 5, reference numeral 530).

D. CLAIM 26 - DEPENDENT

The subject matter of claim 26 is directed to the apparatus of claim 25 further comprising a means for alerting a network manager to the location of the logical port and of the physical port (Specification, p. 11, ll. 10-18; and Figure 1, reference numeral 140).

GROUND OF REJECTION TO BE REVIEWED ON APPEAL

The grounds of rejection to review on appeal are as follows:

A. GROUND OF REJECTION 1

Whether claims 5-11, 15, and 18-27 are not anticipated under 35 U.S.C. §102 by *Ricciulli*, Method of Maintaining Lists of Network Characteristics, U.S. Patent 6,973,040 (December 6, 2005).

B. GROUND OF REJECTION 2

Whether the Examiner failed to state a *prima facie* obviousness rejection of claim 16 under 35 U.S.C. §103 over *Ricciulli* in view of *Hunt et al.*, Network Dispatcher: A Connection Router for Scalable Internet Services, IBM Almaden Research Center, San Jose, CA, available at unizh.ch/home/mazzo/reports/www7conf/fullpapers/1899/com1899.htm.

C. GROUND OF REJECTION 3

Whether the Examiner failed to state a *prima facie* obviousness rejection of claim 17 under 35 U.S.C. §103 over *Ricciulli* in view of *Skirmont et al.*, Method and Apparatus for Load Apportionment Among Physical Interfaces in Data Routers, U.S. Patent 6,553,005 (April 22, 2003).

ARGUMENT

A. GROUND OF REJECTION 1 (Claims 5-11, 15, and 18-27)

The Examiner rejected claims 5-11, 15, and 18-27 under 35 U.S.C. §102 as anticipated by *Ricciulli*. This rejection is manifestly incorrect, as shown below.

A.1. Claims 5-11, 15, 18-20, 24, and 27

A.1.i. Response to Rejection

Claim 5 is a representative claim in this grouping of claims. Claim 5 is as follows:

5. A computer-implemented method of identifying the entry point of an attack upon a device protected by an intrusion detection system, the method comprising the steps of:
 - obtaining intrusion information, from an intrusion detection system, regarding an attack upon a device protected by the intrusion detection system;
 - obtaining network information, from network equipment connected to the device, regarding the attack;
 - determining a logical entry point of the attack using a correlation engine to correlate the intrusion information and the network information;
 - and
 - identifying a physical entry point associated with the logical entry point.

Applicants first address the base rejection of claim 5. Applicants then rebut the Examiner's assertions made in the response to argument section of the final office action of June 27, 2006. In rejecting claim 5, the Examiner states that:

Ricciulli discloses a computer-implemented method of identifying the entry point of an attack upon a device protected by an intrusion detection system, the method comprising the steps of:

Obtaining intrusion information, from an intrusion detection system, regarding an attack upon a device protected by the intrusion detection system (Column 3, lines 16-33);

Obtaining network information, from network equipment connected to the device, regarding the attack (Column 4, line 45 to Column 5, line 2);

Determining a logical entry point (IP addresses, as well as TCP/UDP ports are logical representation used in combination to identify the entry point)

of the attack using a correlation engine to correlate the intrusion information and the network information (Column 3, lines 16-43; and Column 4, line 45 to Column 5, line 2); and

Identifying a physical entry point (the physical entry point is where the router or node actually connects to the network, on its network interface) associated with the logical entry point (Column 3, lines 34-43).

Final office action of June 27, 2006, p. 5.

Ricciulli does not anticipate claim 5 because *Ricciulli* does not teach the claimed step of determining a logical entry point of the attack using a correlation engine to correlate the intrusion information and the network information. Additionally, *Ricciulli* does not anticipate claim 5 because *Ricciulli* does not teach the claimed step of identifying a physical entry point associated with the logical entry point. The Examiner's assertions to the contrary are manifestly incorrect.

Regarding the step of determining a logical entry point of attack, the Examiner asserts that two portions of *Ricciulli* teach this step. In the first portion cited by the Examiner, *Ricciulli* states:

Various embodiments have routers with one or more lists of top-N more seen or most seen network characteristics, for example, destination addresses, in a small cache. This list can vary with time relatively slowly, for example, on the order of seconds for some embodiments. In some embodiments, the cache can have a number of instances of network characteristics substantially equal to or greater than C/F , where C can be a total aggregate capacity of a router, and F can be a minimum sustained flooding rate to detect. For example, to detect a 1 MB/s flooding on a 1 GB/s router, a cache of 1000 instances may be adequate. In one embodiment, listed instances can include a destination address and an ingress port.

When an attack, such as flooding, is detected, a message can be sent upstream by the attacked network node. For example, the message payload can contain a return address R, a network/host address H and/or a cookie generated by the attacked network node.

In some embodiments, an upstream router can look up H to check for a match in one or more lists of the local cache. If a match results, the router can forward a message upstream to appropriate interfaces. This can repeat recursively with routers further upstream.

In some embodiments, if an upstream router does not find H in the local cache, a report message can be sent to R with, for example, interface information of a downstream neighbor and the cookie. In some embodiments, this can be identified as an entry point of the attack, such as the flood.

Ricciulli, col. 3, ll. 16-43 (emphasis supplied).

This portion of *Ricciulli* manifestly does not teach the step of determining a logical entry point of an attack using a correlation engine to correlate the intrusion information and the network information. Instead, this portion of *Ricciulli* indicates that a data packet can be sent from the target of the attack upstream by the attacked network node. The data packet includes a network host address generated by the attacked network node. If an upstream router does not find the host address in the local cache, then a report is sent back to the target of the attack. Thus, the *router* that did not have a host address in its cache is identified as the source of the attack.

Ricciulli does not provide any teaching regarding logical entry points of attack, as claimed. *Ricciulli* does not provide any teaching regarding using a correlation engine to correlate intrusion information and network information to determine a logical entry point of an attack, as claimed. Instead, *Ricciulli* teaches finding a physical entry point of the attack – the router lacking a host address in a cache – by sending a data packet from the target to that router. Nowhere does *Ricciulli* teach that a logical point of attack is determined using a correlation engine as claimed in claim 5. *Ricciulli* does not even discuss logical entry points, except in the context of UDP ports and TCP ports being the type of information that can be compared in *Ricciulli*'s lists in order to identify which physical router is the point of attack.

The Examiner appears to assert that *Ricciulli* does teach something regarding logical entry points of attack. However, the Examiner's assertions regarding IP addresses being logical "representations" is wholly irrelevant to the claimed invention as recited in claim 5.

Specifically, the Examiner states that "IP addresses, as well as TCP/UDP ports are logical representations used in combination to identify the entry point." As a first matter, the Examiner has mischaracterized the claim language. The claim language requires "determining a *logical entry point* of the attack using a *correlation engine* to correlate the *intrusion information and the network*

information.” Whether or not IP addresses and TCP/UDP ports are logical entry points is irrelevant. What is relevant is determining a logical entry point *of an attack*. Nowhere does Ricciulli teach determining a logical entry point *of an attack*. As shown above, Ricciulli teaches finding a *physical router* that is the source of an attack – not a logical entry point of an attack.

Nevertheless, the Examiner asserts that the following text, in combination with the previously cited text, teaches the determining step as claimed:

FIG. 3 shows a flowchart 300 of an aspect of some embodiments for maintaining one or more lists of one or more network characteristics. Various embodiments can alter, add to, delete from, and/or reorder elements of the flowchart 300. In 310, messages can be prevented from transiting the first network node. One embodiment prevents by filtering. Some embodiments prevent, responsive to receiving a message from an attacked network node. The attacked network node may have received a flooding attack and/or a denial of service attack. Such messages can have suspicious instances of network characteristics of lists. The suspicious instances can be associated with attacks on attacked network nodes. In 315, *suspicious instances can be compared with repeatedly updated lists*. If the compare fails to result in a match, prevention can be halted. One embodiment of halting the preventing can include removing the filter. In 320, lists can be repeatedly updated. Instances having low frequency of occurrences can be removed from lists. In various embodiments, the updating can occur at the second network node 140 and/or a third network node.

There are many possible network characteristics that can be matched in 3150. For example, IP source addresses 330, destination IP addresses 335, source TCP ports 340, source UDP ports 345, destination TCP ports 350, destination UDP ports 355, TCP flags 360, and/or ICMP flags 365.

Ricciulli, col. 4, l. 45 through col. 5, l. 2 (emphasis supplied).

This portion of Ricciulli does not teach the step of determining a logical entry point of attack, either alone or together with the previously cited text. This portion of Ricciulli teaches comparing suspicious data packets with repeatedly updated lists. If some aspect of the data packet matches one or more aspects contained in the list, then a suspicious data packet is confirmed to be a data packet associated with an attack. In this case, the source or host of the attack is prevented from sending further data packets to the server. If a suspicious data packet is not associated with an attack, then communication from the source or host is allowed.

However, *Ricciulli* never teaches determining a logical entry point of an attack using a correlation engine, as claimed in claim 5. Although information in data packets is compared to information contained in lists, no comparison is made to determine a *logical entry point of the attack*. Instead, the comparison is made to determine whether a host or source should be blocked from sending further data packets. As described above, the host or source is detected by sending a search data packet upstream in the network to determine which router does not have a host address in a cache. That *physical* router is the source of the attack. Thus, *Ricciulli* does not teach the determining a logical entry point of attack, as claimed.

Additionally, *Ricciulli* does not teach “identifying a physical entry point associated with the logical entry point,” as claimed in claim 5. The Examiner asserts otherwise, citing the following portion of *Ricciulli* for support:

In some embodiments, an upstream router can look up H to check for a match in one or more lists of the local cache. If a match results, the router can forward a message upstream to appropriate interfaces. This can repeat recursively with routers further upstream.

In some embodiments, if an upstream router does not find H in the local cache, a report message can be sent to R with, for example, interface information of a downstream neighbor and the cookie. In some embodiments, this can be identified as an entry point of the attack, such as the flood.

Ricciulli, col. 3, ll. 34-43.

Again, the above-cited text teaches identifying a *physical* router as the source of attack. *Ricciulli* does not mention identifying the logical entry point of attack and does not actually teach identifying a physical entry point associated with the logical entry point of the attack.

Nevertheless, the Examiner also states that, “the physical entry point is where the router or node actually connects to the network, on its network interface.” The Examiner’s statement is irrelevant, whether or not the statement is correct. As *Ricciulli* points out, a *physical* router is identified as the source of an attack. At no point does *Ricciulli* identify a *logical* entry point of attack associated with the *physical* entry point.

As shown above, *Ricciulli* does not teach either the “determining” step or the “identifying” step as claimed in claim 5. Accordingly, *Ricciulli* does not anticipate claim 5.

A.1.ii. Rebuttal to Examiner's Response

In response to the previous response to office action, the Examiner states that:

A network node will detect an attack, thereby obtaining the IP address of an attacking host. This node will then send data to another router, the data including a return address, the attacking host's address, a cookie, a certificate, etc. The packet sent upstream to the current router contains a return IP address for the next downstream router. The current router will then check if the attacking host's information (address, ports, etc.) is within a table. If this information is not in the table, the current router will send a report message to the original network node, said report message comprising interface information of the downstream router, as well as the cookie. It is clear from this that the current router is sent the IP address of the downstream router, correlates network and intrusion information to determine whether the current router has seen the pertinent traffic and, if not, the current router determines that the downstream router's IP address (along with other logical information, such as a logical port) is a logical entry point and creates a message including the physical entry point (interface on the downstream router) to send to the original network node.

Final office action of June 27, 2006, pp. 2-3.

The Examiner's response belies the Examiner's misunderstanding of *Ricciulli* and of the claims. The Examiner's fundamental error is in the Examiner's assertion that that *Ricciulli* determines a logical entry point of attack, as shown above and as proven again below. However, for the sake of Argument, even if the Examiner's invalid assertions were correct, *Ricciulli* still fails to anticipate claim 5.

In particular, the Applicants assume, *arguendo*, that the following statement by the Examiner regarding *Ricciulli* is correct, "the current router determines that the downstream router's IP address (along with other logical information, such as a logical port) is a logical entry point." Thus, Applicants assume, *arguendo*, that an IP address is a logical entry point.

The reason why the Examiner's arguments fall apart is that *Ricciulli* "detects" the IP address (or UDP port or TCP port information) in packets passing through routers and then compares that information to information contained in the lists in the router to determine if the router is the source of the attack. For this reason, *Ricciulli* does not determine a **logical entry point of attack using a correlation engine to correlate the intrusion information and the network information**, as in claim 5. At most, again *arguendo*, *Ricciulli* teaches using a correlation engine to correlate the logical entry point (the IP address) and other information **to determine the physical**

router as the point of attack. Therefore, assuming *arguendo* that the IP address is a logical entry point, *Ricciulli* “detects” the logical entry point of attack and then uses that information elsewhere. *Ricciulli* does not determine a **logical entry point** using a **correlation engine**, as claimed.

For example, see claims 14 and 15 of *Ricciulli* as examples of this technique:

1. A method of maintaining one or more lists of one or more network characteristics of a plurality of messages traveling near at least a first network node coupled to at least a first packet network, comprising:
detecting the plurality of messages traveling near at least the first network node coupled to at least the first packet network, *wherein each of the plurality of messages comprises one or more **network characteristics***; and
updating the one or more lists of the one or more network characteristics of the plurality of messages, such that the one or more lists comprise instances of the one or more network characteristics based on at least a frequency of occurrences of the instances.

...

13. The method of claim 1, *wherein the one or more **network characteristics** comprise one or more source ports*.

14. The method of claim 13, *wherein the one or more source ports comprise one or more **Transmission Control Protocol ports***.

15. The method of claim 13, *wherein the one or more source ports comprise one or more **User Datagram Protocol ports***.

Ricciulli, col. 5, ll. 39-52 and col. 6, ll. 13-20 (emphasis supplied).

Thus, *Ricciulli* simply “picks out” the network characteristic from an information packet. The network characteristic could be a logical entry point in the form of a TCP port or UDP port. *Ricciulli* uses this information to compare to maintained lists, as recited in *Ricciulli*’s claim 1. The rest of *Ricciulli*’s disclosure, as quoted above, is consistent with *Ricciulli*’s claim 1. This mechanism of comparing logical entry points associated with packets passing through routers is not the same as the determination of a logical entry point *of an attack using a correlation engine*, as in claim 5. Therefore, *Ricciulli* does not anticipate claim 5 even if the Examiner’s assertions regarding the nature of logical entry points were correct.

However, the Examiner's assertion that *Ricciulli* teaches determination of a logical entry point of attack is manifestly incorrect. *Ricciulli* determines the *physical entry point* of an attack as follows:

When an attack, such as flooding, is detected, a message can be sent upstream by the attacked network node. For example, the message payload can contain a return address R, a network/host address H and/or a cookie generated by the attacked network node.

In some embodiments, an upstream router can look up H to check for a match in one or more lists of the local cache. If a match results, the router can forward a message upstream to appropriate interfaces. This can repeat recursively with routers further upstream.

In some embodiments, if an upstream router does not find H in the local cache, a report message can be sent to R with, for example, interface information of a downstream neighbor and the cookie. In some embodiments, this can be identified as an entry point of the attack, such as the flood.

Ricciulli, col. 3, ll. 29-43 (quoted above, but repeated here for ease of reference).

Again, *Ricciulli* teaches that a network/host address identified in an attack packet can be compared to a list of addresses in a router. If a match exists, then a message is sent to the next router upstream, whereupon the procedure repeats. Ultimately, when a final router does not find the host address in its local cache, the final router sends a message to the return address in order to identify the final router as the entry point of the attack.

Therefore, *Ricciulli* compares addresses to identify a *physical router* as a source of attack. The Examiner's assertion that *Ricciulli* determines a UDP port address or TCP port address as the source of the attack is incorrect. Instead, *Ricciulli* uses UDP port addresses or TCP port addresses identified in packets passing through routers *as information to compare to the lists contained in the physical router to identify the physical router that is under attack*. Thus, *Ricciulli* does not determine a logical entry point of an attack using a correlation engine to correlate the intrusion information and the network information, as claimed.

Additionally, the Examiner's assertion that *Ricciulli* determines the logical entry point of an attack is incorrect. The Examiner misunderstands what a logical entry point is. A logical entry point is a virtual "port" maintained by a computer's operating system. The virtual port is used to access a network. One of ordinary skill would interpret the claimed term "logical entry point" in

this light. For example, the Board is invited to review the following article provided by a major networking security company, Symantec, which produces Norton Antivirus® software: In that article, Symantec describes a logical entry point as follows:

Port: Logical entry point of a network to your operating system. The operating system has 65,535 logical entry points that can be used by applications to communicate with the outside. Some are “opened” when requested during an outgoing connection, for example, whereas others can remain open permanently to accept connections coming from the outside.

symantec.com/en/aa/home_homeoffice/library/article.jsp?aid=evaluating_security(emphasis in original).

Thus, an Internet address is not a logical entry point, as asserted by the Examiner. An Internet address is just a means of identifying a computer, in much the same way a street address identifies a house.

However, UDP ports or TCP/IP ports are logical entry points. UDP ports and TCP/IP ports are ports maintained by an operating system as logical entry points to the operating system. See also the following article published by Apple Corporation:

You may not be aware of IP ports very often, but you probably use them every day. Servers often deliver more than one type of service, so using the Internet address (URL) of a server is not enough -- you must also tell the server what you want. These requests are made by "port" number. Web service (HTTP) is commonly delivered on port 80, for example. Web browsers are programmed to assume that you want port 80 when you type a URL, such as "www.apple.com". That is why you do not need to be aware of which port you are using for most Web browsing.

Sometimes you need to type a port number when connecting to a service. Whether you must depends on what client software you are using, the service, and how the server is set up.

...

IP stands for "Internet protocol," which can be subdivided into port types such as TCP and UDP. For more on these ports, see "Well-Known" TCP and UDP ports used by Apple software products.

docs.info.apple.com/article.html?artnum=106770.

Although *Ricciulli* does mention UDP ports and TCP ports, *Ricciulli* does so only in the context of identifying types of information that can be pulled from packets passing through the router and then used to compare to information stored in *Ricciulli*'s lists. For example, *Ricciulli* provides that:

There are many possible network characteristics that can be updated in 220. For example, IP source addresses 230, destination IP addresses 235, source TCP ports 240, source UDP ports 245, destination TCP ports 250, destination UDP ports 255, TCP flags 260, and/or ICMP flags 265.

Ricciulli, col. 4, ll. 40-44.

Thus, although *Ricciulli* does identify and compare UDP port and TCP port information in individual packets, *Ricciulli does not identify which of the UDP ports and TCP ports are the logical entry points of attack*, as in claim 5. Certainly, *Ricciulli* does not determine the logical entry point of attack (UDP port or TCP port) *using a correlation engine*, as claimed, because *Ricciulli* draws UDP or TCP port information only from packets passing through routers. Additionally, *Ricciulli* does not correlate the logical entry point of attack to a physical point, as claimed, because *Ricciulli* does not determine the logical entry point of attack in the first place.

Therefore, *Ricciulli* does not teach all of the features of claim 5. Accordingly, *Ricciulli* does not anticipate claim 5.

Applicants now address the Examiner's next argument. The Examiner states that:

It is clear from the above that *Ricciulli* teaches determining a logical entry point of an attack using a correlation engine to correlate the intrusion information with the network information. However, in the sake of clarity, another embodiment of *Ricciulli* discloses that the current router will receive a message from a downstream router, as stated above. The current router will then correlate the network information and intrusion information via tables. If it is found that the current router has been receiving attack traffic, by finding logical information pertaining to the attack within the tables, such as IP source addresses, destination IP addresses, source TCP ports, source UDP ports, destination TCP ports, and destination UDP ports, it will attempt to send the message upstream to another router. If it is determined that the upstream router does not implement the system, the current router will identify itself as the physical entry point and send a message indicating such to the original network node. The logical entry point can be one of many, such as the source IP address, source IP address/source TCP port combo, or any other logical point of entrance through which packets corresponding to an attack travel. Once this logical entry point has been determined by the correlation step,

the current router will identify itself as the physical entry point associated with the logical entry point (or logical entry points, since there could be more than one, as explained above).

Final office action of June 27, 2006, p. 3.

Again, the Examiner's incorrect statements are based on a flawed reading of *Ricciulli*. As shown above, *Ricciulli* does not determine the logical entry point of an attack at all. In another example, in reference to *Ricciulli*'s Figure 2, *Ricciulli* states that:

In 210, messages traveling a first packet network can be detected. The messages can have network characteristics. In 220, lists of network characteristics of messages can be updated, so that lists have instances of network characteristics based on frequency of occurrences of instances. The lists can have a group of more or most frequently occurring instances of the network characteristics. The frequency of occurrences can include a number of occurrences in an amount of time. In one embodiment, a number of instances in each of the lists can be substantially equal to or greater than a quotient. The quotient can include a capacity rate of a router (for example, first network node 110) divided by a threshold flooding rate.

Ricciulli, col. 4, ll. 23-39.

Thus, *Ricciulli* detects *network characteristics* in packets passing through or near routers. The network characteristics can be logical entry points, such as UDP ports or TCP ports. However, *Ricciulli* directly detects these network characteristics and does not “determine a logical entry point of the attack *using a correlation engine* to correlate the intrusion information and the network information.” Instead, *Ricciulli* uses the UDP and TCP port information contained in packets to determine the *physical* router that is the physical entry point of an attack.

Applicants now address the Examiner's next argument. The Examiner states that:

Applicant also argues that the Examiner indicate what IP addresses and TCP/UDP ports are logical representations of. Applicant further argues that using TCP/UDP ports are logical entry points is manifestly incorrect. An IP address is a logical representation of an address for a machine. TCP/UDP ports are logical representations for ports on a machine. Throughout the specification, applicant discusses using logical ports as a possible entry point of an attack, thus the logical ports used as logical entry point within *Ricciulli* is manifestly correct.

Final office action of June 27, 2006, pp. 3-4.

Again, as shown above, *Ricciulli* does not determine which TCP ports and UDP ports are the entry points of attack. Instead, *Ricciulli* uses information regarding TCP ports and UDP ports contained in packets to determine the physical router being attacked.

Applicants now address the Examiner's final argument. The Examiner states that:

Applicant also argues that *Ricciulli* does not teach alerting a network manager to the location of the logical port and of the physical port. Since it has been described above how logical and physical addresses, machines, and ports can be entry points of an attack, and the cited section discloses notifying the relevant ISP or authorities (each being a network manager) about the attack, it is clear to see that the relevant information regarding the attack will be sent to the network managers. Also, as is clear from the specification, alerting a network manager comprises sending a message to a computer or system, so this manager need not be a person. Indeed, the specification does not teach alerting a human manager, only a management center (which is a computer/system). Since this is the case, other sections pertain to this claim as well, such as the sending of a report packet back to the original network node identifying the machine (location of the logical port), interface (physical port), and other information.

Final office action of June 27, 2006, p. 4.

Again, *Ricciulli* does not determine the logical entry point of an attack. Therefore, *Ricciulli* does not teach alerting a network manager to the location of the logical port of attack.

Applicants have successfully rebutted all of the Examiner's arguments. As shown above, *Ricciulli* does not determine the logical entry point of an attack, as in claim 5. *Ricciulli* does not determine the logical entry point of an attack *using a correlation engine*, as in claim 5. Therefore, *Ricciulli* does not teach all of the features of claim 5. Accordingly, *Ricciulli* does not anticipate claim 5 or any other claim in this grouping of claims.

A.2. Claims 21-23, 25, and 26

Claim 21 is a representative claim in this grouping of claims. Claim 21 is as follows:

21. A method of identifying the entry point of an attack upon a device protected by an intrusion detection system, said device one of a plurality of devices connected by a network, the method comprising the computer-implemented steps of:
 - detecting an attack on the device;
 - notifying a correlation engine of the attack on the device;
 - obtaining intrusion information regarding the attack;
 - obtaining network information regarding the attack;

using the correlation engine, correlating the intrusion information and the network information to produce correlation information;
using the correlation information, finding on the network a logical port of connection used by the attack; and
mapping the logical port on the network to a physical port on the network using the correlation engine.

Regarding claim 21, the Examiner states that:

Ricciulli discloses a method of identifying the entry point of an attack upon a device protected by an intrusion detection system, the device being one of a plurality of devices connected by a network, the method comprising the computer-implemented steps of:

Detecting an attack on the device (Column 3, lines 16-33);

Notifying a correlation engine of the attack on the device (Column 3, lines 16-33);

Obtaining intrusion information regarding the attack (Column 3, lines 16-33);

Obtaining network information regarding the attack (Column 4, line 45 to Column 5, line 2);

Using the correlation engine, correlating the intrusion information and the network information to produce correlation information (Column 3, lines 16-43; and Column 4, line 45 to Column 5, line 2);

Using the correlation information, finding on the network a logical port of connection used by the attack (Column 3, lines 16-43; and Column 4, line 45 to Column 5, line 2); and

Mapping the logical port on the network to a physical port on the network using the correlation engine (Column 3, lines 34-43).

Final office action of June 27, 2006, pp. 7-8.

As shown above *vis-à-vis* the response to the rejection of claim 5, *Ricciulli* does not teach the claimed feature of, “using the correlation information, finding on the network a logical port of connection used by the attack,” as in claim 21. Instead, *Ricciulli* uses logical port information (UDP port and TCP port) contained in packets passing through a router to identify which physical router is the source of an attack. In any case, *Ricciulli* does not teach finding the logical port of connection used by the attack using correlation information, as in claim 21.

Additionally, *Ricciulli* does not teach, “mapping the logical port on the network to a physical port on the network using the correlation engine,” as in claim 21. The Examiner asserts otherwise, referring to the following portion of *Ricciulli*:

In some embodiments, an upstream router can look up H to check for a match in one or more lists of the local cache. If a match results, the router can forward a message upstream to appropriate interfaces. This can repeat recursively with routers further upstream.

In some embodiments, if an upstream router does not find H in the local cache, a report message can be sent to R with, for example, interface information of a downstream neighbor and the cookie. In some embodiments, this can be identified as an entry point of the attack, such as the flood.

Ricciulli, col. 3, ll. 34-43.

In this text, the term “H” refers to the host address and the term “R” refers to the return address. See, for example, the immediately preceding paragraph of *Ricciulli*.

When an attack, such as flooding, is detected, a message can be sent upstream by the attacked network node. For example, the message payload can contain a return address R, a network/host address H and/or a cookie generated by the attacked network node.

Ricciulli, col. 3, ll. 29-33.

Thus, *Ricciulli* teaches that an upstream router can look up the host address to check for a match in the lists. If a match results, the router forwards a message upstream to another router upstream. If a match does not result, then a report message is sent to the return address. The *upstream router* is identified as the entry point of the attack.

Ricciulli’s exact statement is that, “this can be identified as an entry point of the attack.” *Ricciulli* does not identify the term to which the term “this” refers. However, in the context of the surrounding text, surrounding paragraphs, and *Ricciulli*’s claims, one of ordinary skill would realize that the upstream router is identified as the source of the attack. The term “this” in *Ricciulli*’s sentence could not refer to any other term in the cited paragraph, as the result would not make sense.

In any case, *Ricciulli* does not teach, “mapping the logical port on the network to a physical port on the network using the correlation engine,” as in claim 21. *Ricciulli* does not map a logical port to a physical port. Rather, in the cited text *Ricciulli* teaches correlating information in packets

passing through routers to information in lists contained in routers to identify the router under attack. Although *Ricciulli* does use logical port address information in the packets to identify a physical port, *Ricciulli* never teaches *mapping* the logical ports to a physical port, as claimed. These two techniques are completely different from each other.

Therefore, *Ricciulli* does not teach all of the features of claim 21. Accordingly, *Ricciulli* does not anticipate claim 21 or any other claim in this grouping of claims.

B. GROUND OF REJECTION 2 (Claim 16)

Claim 16 is as follows:

16. The computer-implemented method of claim 5, wherein the network equipment includes a network dispatcher.

Regarding claim 16, the Examiner states that:

Ricciulli does not disclose that the network equipment includes a network dispatcher.

ND, however, discloses that the network equipment includes a network dispatcher (Pages 1-2, Introduction, Paragraphs 1-4). It would have been obvious to one of ordinary skill in the art at the time of applicant's invention to incorporate the network dispatcher of ND into the intrusion detection system of *Ricciulli* in order to allow the system to protect a broader range of network equipment, thus increasing the types of routers that can be used and protected by the system, and to reach those customers that use network dispatchers.

Final office action of June 27, 2006, pp. 10-11.

B.1. The Proposed Combination Does Not Teach or Suggest All of the Features of Claim 16

The Examiner bears the burden of establishing a *prima facie* case of obviousness based on prior art when rejecting claims under 35 U.S.C. §103. *In re Fritch*, 972 F.2d 1260, 23 U.S.P.Q.2d 1780 (Fed. Cir. 1992). A *prima facie* case of obviousness is established when the teachings of the prior art itself suggest the claimed subject matter to a person of ordinary skill in the art. *In re Bell*, 991 F.2d 781, 783, 26 U.S.P.Q.2d 1529, 1531 (Fed. Cir. 1993). All limitations of the claimed invention must be considered when determining patentability. *In re Lowry*, 32 F.3d 1579, 1582, 32

U.S.P.Q.2d 1031, 1034 (Fed. Cir. 1994). For an invention to be *prima facie* obvious, the prior art must teach or suggest all claim limitations. *In re Royka*, 490 F.2d 981, 180 USPQ 580 (CCPA 1974). In the case at hand, the cited references do not teach or suggest all of the features of the claims, arranged as they are in the claims.

The Examiner has failed to state a *prima facie* obviousness rejection because the proposed combination, when considered as a whole, does not teach or suggest all of the features of claim 16. Claim 16 depends from claim 5. As shown above, *Ricciulli* does not teach all of the features of claim 5. *Ricciulli* is devoid of disclosure regarding determination of logical entry points of attack and is devoid of disclosure regarding use of correlation engines to determine logical entry points of attack. Therefore, *Ricciulli* does not suggest these claimed features.

Hunt also does not teach the claimed features not shown in *Ricciulli*. *Hunt* is directed towards the design of a network dispatcher. *Hunt* is devoid of disclosure regarding determination of logical entry points of attack and is devoid of disclosure regarding use of correlation engines to determine logical entry points of attack. Therefore, *Hunt* also does not suggest these claimed features.

Because both *Ricciulli* and *Hunt* do not teach or suggest the features of claim 16, the proposed combination when considered as a whole does not teach or suggest the features of claim 16. Accordingly, under the standards of *In re Lowry*, the Examiner has failed to state a *prima facie* obviousness rejection.

B.2. The Examiner Failed To State a Proper Teaching, Suggestion, or Motivation To Combine the References

Regarding a teaching, suggestion, or motivation to combine the references, the Examiner states that:

It would have been obvious to one of ordinary skill in the art at the time of applicant's invention to incorporate the network dispatcher of ND into the intrusion detection system of *Ricciulli* in order to allow the system to protect a broader range of network equipment, thus increasing the types of routers that can be used and protected by the system, and to reach those customers that use network dispatchers.

Final office action of June 27, 2006, p. 11.

A proper *prima facie* case of obviousness cannot be established by combining the teachings of the prior art absent some teaching, incentive, or suggestion supporting the combination. *In re Napier*, 55 F.3d 610, 613, 34 U.S.P.Q.2d 1782, 1784 (Fed. Cir. 1995); *In re Bond*, 910 F.2d 831, 834, 15 U.S.P.Q.2d 1566, 1568 (Fed. Cir. 1990). In the case at hand the Examiner has failed to establish a proper teaching, incentive, or suggestion supporting the combination and no such teaching, incentive, or suggestion exists.

The Examiner has not stated a motivation to combine the references. Instead, the Examiner has cited a purported advantage of combining the references. An advantage is not necessarily a motivation. To serve as a motivation to combine the references to achieve the claimed invention, one of ordinary skill must, logically, both recognize the advantage and have a reason to implement the advantage.

In the case at hand, the Examiner has provided no reason why one of ordinary skill would recognize the proposed advantage. The Examiner has provided no reason why one of ordinary skill would, without a-priori knowledge, think to incorporate network dispatchers in the rest of the claimed invention. Instead, the Examiner has only picked a feature known in the art and then combined that feature with the base claim without identifying why one of ordinary skill would combine the references in the first place.

Therefore, the Examiner has failed to provide a proper teaching, suggestion or motivation to combine the references. Accordingly, the Examiner has failed to state a *prima facie* obviousness rejection against claim 16.

B.3. No Motivation Exists to Combine *Ricciulli* and *Hunt* Because They Address Different Problems

One of ordinary skill would not combine the references to achieve the invention of claim 16 because the references are directed towards solving different problems. It is necessary to consider the reality of the circumstances--in other words, common sense--in deciding in which fields a person of ordinary skill would reasonably be expected to look for a solution to the problem facing the inventor. *In re Oetiker*, 977 F.2d 1443 (Fed. Cir. 1992); *In re Wood*, 599 F.2d 1032, 1036, 202 U.S.P.Q. 171, 174 (CCPA 1979). In the case at hand, the cited references address distinct problems. Thus, no common sense reason exists to establish that one of ordinary skill would

reasonably be expected to look for a solution to the problem facing the inventor. Accordingly, no teaching, suggestion, or motivation exists to combine the references and the Examiner has failed to state a *prima facie* obviousness rejection of claim 16.

For example, *Ricciulli* is directed to solving the problem of detecting the router that is the physical entry point of a denial of service attack. For example, *Ricciulli* provides that:

Potential loss of revenue caused by preempting reliable TCP communications is enormous, and therefore adequate mechanisms for dealing with SYN flooding are needed. Current SYN flooding defense mechanisms seem to have greatly mitigated the problem by making it harder for an attacker to negatively affect service. The most popular approach uses a "brute force" technique. In this approach, the TCP "connection pending" data structure (implementing the connection request queue) is made sufficiently large that an average attacker, to be successful, would need to flood connection requests at a rate exceeding reasonable bandwidth capabilities. This solution, although sometimes very practical, requires large amounts of protected kernel memory and may slow down the server response time for looking up connections in the vast "connection pending" data structure. Other less popular techniques use one-way hash functions (with Internet "cookies") to verify the authenticity of connection requests and therefore eliminate unnecessary memory allocation. Some of these latter techniques can introduce changes in the TCP signaling behavior and are therefore less favored. Firewall approaches actively monitor the TCP signaling traffic to detect possible attacks and inject ad-hoc signaling messages in the network to mitigate the denial-of-service attack. These approaches are awkward because they introduce additional administrative complexity, may introduce significant delays for legitimate connection establishment, or may expose the system to different, though arguably less severe, kinds of vulnerabilities.

No one mechanism seems to provide an optimal solution, and thus a careful protection approach is usually constructed by using a combination of techniques. What is needed is a solution that can complement or replace existing solutions.

Ricciulli, col. 2, ll. 4-26.

On the other hand, *Hunt* is directed to the problem of designing network dispatchers. For example, *Hunt* provides as follows:

Network Dispatcher (ND) is a TCP connection router that supports load sharing across several TCP servers. Prototypes of Network Dispatcher were used to support several large scale high-load Web sites. Network Dispatcher provides a fast IP packet-forwarding kernel-extension to the

TCP/IP stack. Load sharing is supported by a user-level manager process that monitors the load on the servers and controls the connection allocation algorithm in the kernel extension. This paper describes the design of Network Dispatcher, outlines Network Dispatcher's performance in the context of http traffic, and presents several of its features including high-availability, support for WANs, and client affinity.

Hunt, Abstract (emphasis in original).

Based on the plain disclosures of the references themselves, the references address completely distinct problems that are unrelated to each other. The problem of detecting the router that is the physical entry point of a denial of service attack is completely distinct from the problem of designing network dispatchers.

Because the references address completely distinct problems, one of ordinary skill would have no reason to combine or otherwise modify the references to achieve the invention of claim 16. Thus, no proper teaching, suggestion, or motivation exists to combine the references in the manner suggested by the Examiner. Accordingly, the Examiner has failed to state a *prima facie* obviousness rejection against claim 16.

B.4. The Examiner Used Impermissible Hindsight When Combining the References

The Examiner failed to state a *prima facie* obviousness rejection because the Examiner used impermissible hindsight when fashioning the rejections. "It is impermissible within the framework of section 103 to pick and choose from any one reference only so much of it as will support a given position, to the exclusion of other parts necessary to the full appreciation of what such reference fairly suggests to one of ordinary skill in the art." *In re Hedges*, 228 U.S.P.Q. 685, 687 (Fed. Cir. 1986). Additionally, personal opinion cannot be substituted for what the prior art teaches because a *prima facie* case of obviousness is established when the teachings of the prior art itself suggest the claimed subject matter to a person of ordinary skill in the art. *In re Bell*, 991 F.2d 781, 783, 26 U.S.P.Q.2d 1529, 1531 (Fed. Cir. 1993). In viewing the references as a whole, one of ordinary skill would look to the problems addressed by the references in determining whether to combine the references.

As shown above, the Examiner provided no reason why one of ordinary skill would recognize the proposed advantage. The Examiner has provided no reason why one of ordinary skill

would, without a-priori knowledge, think to incorporate network dispatchers in the rest of the claimed invention. Instead, the Examiner only provided a hypothetical advantage to combining the references.

Given the disparity between the references, the Examiner clearly searched the prior art for the term “network dispatcher,” found a reference, and then tried to glue the references together to achieve the invention of claim 16 without regard to the fact that one of ordinary skill would have no reason to look to combine the references in the first place. The Examiner may not “pick and choose” features in the prior art and then combine them together in this manner. The Examiner’s attempt to do so is impermissible hindsight under the standards of *In re Hedges*. Accordingly, the Examiner has failed to state a *prima facie* obviousness rejection against claim 16.

C. GROUND OF REJECTION 3 (Claim 17)

Claim 17 is as follows:

17. The computer-implemented method of claim 5, wherein the network equipment includes a load balancer.

Regarding claim 17, the Examiner states that:

Ricciulli does not disclose that the network equipment includes a load balancer.

Skirmont, however, discloses that the network equipment includes a load balancer (Column 5, lines 52-67). It would have been obvious to one of ordinary skill in the art at the time of applicant’s invention to incorporate the load balancing system of Skirmont in the intrusion detection system of Ricciulli in order to map packets that have a common source and destination by strict physical paths, while at the same time accomplishing efficient load balancing along the same physical paths, thus protecting packets being received out of order, and consequently being lost/discarded (Column 1, lines 41-64; and Column 2, lines 20-50).

Final office action of June 27, 2006, p. 11.

C.1. The Proposed Combination Does Not Teach or Suggest All of the Features of Claim 17

The Examiner bears the burden of establishing a *prima facie* case of obviousness based on prior art when rejecting claims under 35 U.S.C. § 103. *In re Fritch*, 972 F.2d 1260, 23 U.S.P.Q.2d

1780 (Fed. Cir. 1992). A *prima facie* case of obviousness is established when the teachings of the prior art itself suggest the claimed subject matter to a person of ordinary skill in the art. *In re Bell*, 991 F.2d 781, 783, 26 U.S.P.Q.2d 1529, 1531 (Fed. Cir. 1993). All limitations of the claimed invention must be considered when determining patentability. *In re Lowry*, 32 F.3d 1579, 1582, 32 U.S.P.Q.2d 1031, 1034 (Fed. Cir. 1994). For an invention to be *prima facie* obvious, the prior art must teach or suggest all claim limitations. *In re Royka*, 490 F.2d 981, 180 USPQ 580 (CCPA 1974). In the case at hand, the cited references do not teach or suggest all of the features of the claims, arranged as they are in the claims.

The Examiner has failed to state a *prima facie* obviousness rejection because the proposed combination, when considered as a whole, does not teach or suggest all of the features of claim 17. Claim 17 depends from claim 5. As shown above, *Ricciulli* does not teach all of the features of claim 5. *Ricciulli* is devoid of disclosure regarding determination of logical entry points of attack and is devoid of disclosure regarding use of correlation engines to determine logical entry points of attack. Therefore, *Ricciulli* does not suggest these claimed features.

Skirmont also does not teach the claimed features not shown in *Ricciulli*. *Skirmont* is directed towards load balancing among physical routers. *Skirmont* is devoid of disclosure regarding determination of logical entry points of attack and is devoid of disclosure regarding use of correlation engines to determine logical entry points of attack. Therefore, *Skirmont* also does not suggest these claimed features.

Because both *Ricciulli* and *Skirmont* do not teach or suggest the features of claim 17, the proposed combination when considered as a whole does not teach or suggest the features of claim 17. Accordingly, under the standards of *In re Lowry*, the Examiner has failed to state a *prima facie* obviousness rejection.

C.2. The Examiner Failed To State a Proper Teaching, Suggestion, or Motivation To Combine the References

Regarding a teaching, suggestion, or motivation to combine the references, the Examiner states that:

It would have been obvious to one of ordinary skill in the art at the time of applicant's invention to incorporate the load balancing system of *Skirmont* in the intrusion detection system of *Ricciulli* in order to map packets that

have a common source and destination by strict physical paths, while at the same time accomplishing efficient load balancing along the same physical paths, thus protecting packets being received out of order, and consequently being lost/discarded (Column 1, lines 41-64; and Column 2, lines 20-50).

Final office action of June 27, 2006, p. 11.

A proper *prima facie* case of obviousness cannot be established by combining the teachings of the prior art absent some teaching, incentive, or suggestion supporting the combination. *In re Napier*, 55 F.3d 610, 613, 34 U.S.P.Q.2d 1782, 1784 (Fed. Cir. 1995); *In re Bond*, 910 F.2d 831, 834, 15 U.S.P.Q.2d 1566, 1568 (Fed. Cir. 1990). In the case at hand the Examiner has failed to establish a proper teaching, incentive, or suggestion supporting the combination and no such teaching, incentive, or suggestion exists.

The Examiner has not stated a motivation to combine the references. Instead, the Examiner has cited a purported advantage of combining the references. An advantage is not necessarily a motivation. To serve as a motivation to combine the references to achieve the claimed invention, one of ordinary skill must, logically, both recognize the advantage and have a reason to implement the advantage.

In the case at hand, the Examiner has provided no reason why one of ordinary skill would recognize the proposed advantage. The Examiner has provided no reason why one of ordinary skill would, without a-priori knowledge, think to incorporate load balancers in the rest of the claimed invention. Instead, the Examiner has only picked a feature known in the art and then combined that feature with the base claim without identifying why one of ordinary skill would combine the references in the first place.

Therefore, the Examiner has failed to provide a proper teaching, suggestion or motivation to combine the references. Accordingly, the Examiner has failed to state a *prima facie* obviousness rejection against claim 17.

C.3. No Motivation Exists to Combine *Ricciulli* and *Skirmont* Because They Address Different Problems

One of ordinary skill would not combine the references to achieve the invention of claim 17 because the references are directed towards solving different problems. It is necessary to consider

the reality of the circumstances--in other words, common sense--in deciding in which fields a person of ordinary skill would reasonably be expected to look for a solution to the problem facing the inventor. *In re Oetiker*, 977 F.2d 1443 (Fed. Cir. 1992); *In re Wood*, 599 F.2d 1032, 1036, 202 U.S.P.Q. 171, 174 (CCPA 1979). In the case at hand, the cited references address distinct problems. Thus, no common sense reason exists to establish that one of ordinary skill would reasonably be expected to look for a solution to the problem facing the inventor. Accordingly, no teaching, suggestion, or motivation exists to combine the references and the Examiner has failed to state a *prima facie* obviousness rejection of claim 17.

For example, *Ricciulli* is directed to solving the problem of detecting the router that is the physical entry point of a denial of service attack. On the other hand, *Skirmont* is directed to the problem of load apportionment among physical interfaces in data routers. For example, *Skirmont* provides as follows:

In current art when a packet is received at a router the packets headers are read and typically a forwarding table is consulted to determine the next hop for the packet. This next hop table contains, among other things, the identity of the egress interface to be used and how to send the packet internally to that location. A problem in current art is that the egress interface may well be a defined interface comprising several actual physical egress ports. The problem then is one of determining which of the actual physical egress ports to use. One solution is to simply do another software table lookup. This is not difficult for software based routing elements, but is less than ideal for a high-speed hardware based solution where memory space and ASIC pins may well be limited.

What is clearly needed for the new generation of very high-speed and more sophisticated routers is a method and system for mapping IP packets that have common source and destination by strict physical paths, while at the same time accomplishing efficient load balancing along the same physical paths.

Skirmont, col. 2, ll. 6-25.

Based on the plain disclosures of the references themselves, the references address completely distinct problems that are unrelated to each other. The problem of detecting the router that is the physical entry point of a denial of service attack is completely distinct from the problem of load apportionment among physical interfaces in data routers.

Because the references address completely distinct problems, one of ordinary skill would have no reason to combine or otherwise modify the references to achieve the invention of claim 17. Thus, no proper teaching, suggestion, or motivation exists to combine the references in the manner suggested by the Examiner. Accordingly, the Examiner has failed to state a *prima facie* obviousness rejection against claim 17.

C.4. The Examiner Used Impermissible Hindsight When Combining the References

The Examiner failed to state a *prima facie* obviousness rejection because the Examiner used impermissible hindsight when fashioning the rejections. "It is impermissible within the framework of § 103 to pick and choose from any one reference only so much of it as will support a given position, to the exclusion of other parts necessary to the full appreciation of what such reference fairly suggests to one of ordinary skill in the art." *In re Hedges*, 228 U.S.P.Q. 685, 687 (Fed. Cir. 1986). Additionally, personal opinion cannot be substituted for what the prior art teaches because a *prima facie* case of obviousness is established when the teachings of the prior art itself suggest the claimed subject matter to a person of ordinary skill in the art. *In re Bell*, 991 F.2d 781, 783, 26 U.S.P.Q.2d 1529, 1531 (Fed. Cir. 1993). In viewing the references as a whole, one of ordinary skill would look to the problems addressed by the references in determining whether to combine the references.

As shown above, the Examiner provided no reason why one of ordinary skill would recognize the proposed advantage. The Examiner has provided no reason why one of ordinary skill would, without a-priori knowledge, think to load balancers in the rest of the claimed invention. Instead, the Examiner only provided a hypothetical advantage to combining the references.

Given the disparity between the references, the Examiner clearly searched the prior art for the term "load balancer," found a reference, and then tried to glue the references together to achieve the invention of claim 17 without regard to the fact that one of ordinary skill would have no reason to look to combine the references in the first place. The Examiner may not "pick and choose" features in the prior art and then combine them together in this manner. The Examiner's attempt to do so is impermissible hindsight under the standards of *In re Hedges*. Accordingly, the Examiner has failed to state a *prima facie* obviousness rejection against claim 17.

D. CONCLUSION

As shown above, *Ricciulli* does not anticipate the claims. Similarly, the Examiner has failed to state a *prima facie* obviousness rejection against the claims. Therefore, Applicants request that the Board of Patent Appeals and Interferences reverse the rejections. Additionally, Applicants request that the Board direct the Examiner to allow the claims.

/Theodore D. Fay III/

Theodore D. Fay III

Reg. No. 48,504

YEE & ASSOCIATES, P.C.

PO Box 802333

Dallas, TX 75380

(972) 385-8777

CLAIMS APPENDIX

The text of the claims involved in the appeal is as follows:

5. A computer-implemented method of identifying the entry point of an attack upon a device protected by an intrusion detection system, the method comprising the steps of:
 - obtaining intrusion information, from an intrusion detection system, regarding an attack upon a device protected by the intrusion detection system;
 - obtaining network information, from network equipment connected to the device, regarding the attack;
 - determining a logical entry point of the attack using a correlation engine to correlate the intrusion information and the network information; and
 - identifying a physical entry point associated with the logical entry point.
6. The computer-implemented method of claim 5, wherein the intrusion information includes an address.
7. The computer-implemented method of claim 6, wherein the address is a source address.
8. The computer-implemented method of claim 6, wherein the address is a destination address.

9. The computer-implemented method of claim 6, wherein the network information includes a logical port identifier of a logical port associated with the address.
10. The computer-implemented method of claim 9, wherein the step of determining a logical entry point includes the step of finding, in the network information, the logical port identifier of the logical port associated with the address.
11. The computer-implemented method of claim 9, wherein the step of identifying a physical entry point includes the step of identifying a physical port associated with the logical port.
15. The computer-implemented method of claim 5, wherein the network equipment includes a firewall with routing function.
16. The computer-implemented method of claim 5, wherein the network equipment includes a network dispatcher.
17. The computer-implemented method of claim 5, wherein the network equipment includes a load balancer.
18. The computer-implemented method of claim 5, wherein the intrusion detection system includes network based intrusion detection equipment.

19. The computer-implemented method of claim 5, wherein the intrusion detection system includes host based intrusion detection equipment.

20. The computer-implemented method of claim 5, wherein the intrusion detection system includes application based intrusion detection equipment.

21. A method of identifying the entry point of an attack upon a device protected by an intrusion detection system, said device one of a plurality of devices connected by a network, the method comprising the computer-implemented steps of:

detecting an attack on the device;

notifying a correlation engine of the attack on the device;

obtaining intrusion information regarding the attack;

obtaining network information regarding the attack;

using the correlation engine, correlating the intrusion information and the network information to produce correlation information;

using the correlation information, finding on the network a logical port of connection used by the attack; and

mapping the logical port on the network to a physical port on the network using the correlation engine.

22. The method of claim 21 comprising the further step of: alerting a network manager to the location of the logical port and of the physical port.

23. The method of claim 21 wherein the step of mapping is performed using the correlation engine.
24. The method of claim 21 wherein:
the intrusion information includes an address; and
the network information includes a logical port identifier of a logical port associated with the address.
25. An apparatus for detecting a point of an attack on a network, the apparatus comprising:
network equipment for connecting a protected device to a network;
an intrusion detection system comprising intrusion detection equipment;
a correlation engine adapted to:
receive a notification of an attack on the protected device;
receive intrusion information regarding the attack;
receive network information regarding the attack, wherein the network information pertains to the network;
correlate the intrusion information and the network information to produce correlation information;
use the correlation information to find on the network a logical port of connection used by the attack; and
map the logical port on the network to a physical port on the network using the correlation engine.

26. The apparatus of claim 25 further comprising:
means for alerting a network manager to the location of the logical port and of the physical port.
27. The apparatus of claim 25 wherein:
the intrusion information includes an address; and
the network information includes a logical port identifier of a logical port associated with the address.

EVIDENCE APPENDIX

There is no evidence to be presented.

RELATED PROCEEDINGS APPENDIX

There are no related proceedings.